

WE CLAIM:

1. A data storage system comprising:
 - a communication network;
 - a client application coupled to the network and generating an access request for stored data, wherein the client application lacks a priori knowledge of the location of the requested data;
 - an intermediary server coupled to the network to receive the request;
 - one or more data storage devices accessible through the intermediary server and having a plurality of data units stored at selected locations therein;
 - a storage server having knowledge of the location of data units in the storage devices and having an interface for communicating with the intermediary servers;
 - processes within the intermediary server responsive to a received data access request for communicating with the storage server to obtain knowledge about the location of requested data from the data in response to a received client request; and
 - processes within the intermediary server for obtaining the data from the specific location and serving the data to the requesting client application.
2. The system of claim 1 wherein the data is returned such that the client remains unaware of the specific location of the data.
3. The system of claim 1 wherein the intermediary server has a lower latency connection to the client application than does the storage server.
4. The system of claim 1 wherein at least some of the storage devices comprise direct attached storage for the intermediary server.

5. The system of claim 1 wherein at least some of the storage devices comprise network attached storage.

6. The system of claim 1 wherein at least some of the storage device are configured as a storage area network.

7. The system of claim 1 wherein the access request is represented by a token.

8. The system of claim 1 wherein the processes for communicating with the storage server further comprises transmission of a token representing the requested data.

9. The system of claim 1 wherein the processes for communicating with the storage server further comprises processes for receiving a resource locator from the storage server.

10. The system of claim 1 wherein the processes for communicating with the storage server further comprise processes for receiving a file name and file path from the intermediary server.

11. A method for managing on-network data storage comprising the acts of:

providing a communication network;

5 receiving requests for data within an intermediary server from a plurality of external client applications coupled to the network;

storing units of data in one or more data storage devices accessible to the intermediary server;

10 associating each storage request with a token representing the request;

sending the token to a storage server coupled to the network and having an interface for communicating with the intermediary server;

causing the storage server to return specific
15 location information corresponding to the request
associated with the received token;
causing the intermediary server to access the data
storage mechanism using the specific location information
to retrieve data at the specific location; and
20 delivering the retrieved data to the client
application that generated the request.

12. A method for transferring data between network-connected computers comprising the acts of:

storing a data object at a specific location in a
network-connected storage mechanism;

5 transmitting a token representing the data object
from a first network-connected computer to an
intermediary computer;

10 in the intermediary computer, using the token to
identify the specific storage location of the data
object;

causing the storage mechanism to transfer the data
object to a second network-connected computer.

13. The method of claim 12 wherein the step of
sending the token further comprises sending an
identification of the second network-connected computer.

14. The method of claim 12 wherein the act of
transferring the data object comprises transferring the
data object to a plurality of network-connected
computers.

15 The method of claim 12 further comprising:

storing copies of the data object at multiple
network-connected storage mechanisms;

5 using the intermediary computer to select one of the
multiple network-connected storage mechanisms; and

causing the selected network-connected storage mechanism to transfer the data object to a second network-connected computer.

16. The method of claim 12 wherein the step of causing the storage mechanism to transfer the data object to a second network-connected computer comprises:

5 transferring the data object to a front-end server topologically close to the second network-connected computer; and

transferring the data object from the front-end server to the second network-connected computer.

17. The method of claim 12 wherein the data object at the specific location is referred to as a primary data object, the method further comprising:

5 causing the network-connected storage mechanism to proactively redistribute data objects by transferring in addition to the primary data object, one or more data objects that are sequentially related to the primary data object.

18. A data distribution service comprising:

one or more data storage mechanisms holding a plurality of data objects at specific non-public locations;

5 an interface for receiving tokens, the tokens associated with particular ones of the data objects and the tokens lacking specific location information indicating the locations of the data objects in the one or more data storage mechanisms; and

10 in exchange for payment, supplying the specific non-public locations of the data objects associated with the received tokens.

19. A method for version control of a data object comprising:

receiving a token representing a first version of a
data object;
5 using the token to identify second version of the
data object; and
identifying a specific storage location of the
second version data object in response to the received
token.